

## **SPECIFICATION**

(Atty. Docket No. RA158)

### **TO WHOM IT MAY CONCERN:**

Be it known that Rich Perego, a citizen of United States of America, and a resident of 4337 Renaissance Dr. #120, San Jose, CA 95134, Ely Tsem, a citizen of the United States of America and a resident of 684 Riviera Drive, Los Altos, CA 94024, and Stefanos Sidiropoulos a citizen of Greece and a resident of 731 Ellsworth Street Palo Alto, California 94306 have invented a new and useful:

### **Memory System Including a Point-to-Point Linked Memory Subsystem**

of which the following is a specification:

## **BACKGROUND OF THE INVENTION**

This invention relates to memory systems, memory subsystems, memory modules or a system having memory devices. More specifically, this invention is directed toward memory system architectures which may include integrated circuit devices such as one or  
5 more controllers and a plurality of memory devices.

Some contemporary memory system architectures may demonstrate tradeoffs between cost, performance and the ability to upgrade, for example; the total memory capacity of the system. Memory capacity is commonly upgraded via memory modules or cards featuring a connector/socket interface. Often these memory modules are connected  
10 to a bus disposed on a backplane to utilize system resources efficiently. System resources include integrated circuit die area, package pins, signal line traces, connectors, backplane board area, just to name a few. In addition to upgradeability, many of these contemporary memory systems also require high throughput for bandwidth intensive applications, such as graphics.

15 With reference to FIGURE 1, a representational block diagram of a conventional memory system employing memory modules is illustrated. Memory system 100 includes memory controller 110 and modules 120a-120c. Memory controller 110 is coupled to modules 120a-120c via control/address bus 130, data bus 140, and corresponding module control lines 150a-150c. Control/address bus 130 typically comprises a plurality of  
20 address lines and control signals (e.g., RAS, CAS and WE).

The address lines and control signals of control/address bus 130 are bussed and “shared” between each of modules 120a-120c to provide row/column addressing and read/write, precharge, refresh commands, etc., to memory devices on a selected one of

modules 120a-120c. Individual module control lines 150a-150c are typically dedicated to a corresponding one of modules 120a-120c to select which of modules 120a-120c may utilize the control/address bus 130 and data bus 140 in a memory operation.

Here and in the detailed description to follow, “bus” denotes a plurality of signal  
5 lines, each having more than two connection points for “transceiving” (i.e., transmitting or receiving). Each connection point electrically connects or couples to a transceiver (i.e., a transmitter-receiver) or one of a single transmitter or receiver circuit.

With further reference to FIGURE 1, memory system 100 may provide an upgrade path through the usage of modules 120a-120c. A socket and connector interface  
10 may be employed which allows each module to be removed and replaced by a memory module that is faster or includes a higher capacity. Memory system 100 may be configured with unpopulated sockets or less than a full capacity of modules (i.e., empty sockets/connectors) and provided for increased capacity at a later time with memory expansion modules. Since providing a separate group of signals (e.g., address lines and  
15 data lines) to each module is avoided using the bussed approach, system resources in memory system 100 are efficiently utilized.

U.S. Patent 5,513,135 discloses a contemporary dual inline memory module (DIMM) having one or more discrete buffer devices. In this patent, the discrete buffer devices are employed to buffer or register signals between memory devices disposed on  
20 the module and external bussing (such as control/address bus 130 in memory system 100). The discrete buffer devices buffer or register incoming control signals such as RAS, and CAS, etc., and address signals. Local control/address lines are disposed on the contemporary memory module to locally distribute the buffered or registered control and

address signals to each memory device on the module. By way of note, the discrete buffer devices buffer a subset of all of the signals on the memory module since data path signals (e.g., data bus 140 in FIGURE 1) of each memory device are connected directly to the external bus.

5           In addition to the discrete buffer device(s), a phase locked Loop (PLL) device may be disposed on the contemporary DIMM described above. The PLL device receives an external clock and generates a local phase adjusted clock for each memory device as well as the discrete buffer devices.

Modules such as the DIMM example disclosed in U.S. Patent 5,513,135 feature  
10    routed connections between input/outputs (I/Os) of each memory device and connector pads disposed at the edge of the module substrate. These routed connections introduce long stub lines between the signal lines of the bus located off of the module (e.g., control address bus 130 and data bus 140), and memory device I/Os. A stub line is commonly  
15    known as a routed connection that deviates from the primary path of a signal line. Stub lines commonly introduce impedance discontinuities to the signal line. Impedance discontinuities may produce undesirable voltage reflections manifested as signal noise that may ultimately limit system operating frequency.

Examples of contemporary memory systems employing buffered modules are illustrated in FIGURES 2A and 2B. FIGURE 2A illustrates a memory system 200 based  
20    on a Rambus™ channel architecture and FIGURE 2B illustrates a memory system 210 based on a Synchronous Link architecture. Both of these systems feature memory modules having buffer devices 250 disposed along multiple transmit/receive connection points of bus 260. In both of these examples, the lengths of stubs are significantly

shortened in an attempt to minimize signal reflections and enable higher bandwidth characteristics. Ultimately however, memory configurations such as the ones portrayed by memory systems 100, 200 and 210 may be significantly bandwidth limited by the electrical characteristics inherent in the bussed approach as described below.

5           In the bussed approach exemplified in FIGURES 1, 2A and 2B, the signal lines of the bussed signals become loaded with a (load) capacitance associated with each bus connection point. These load capacitances are normally attributed to components of input/output (I/O) structures disposed on an integrated circuit (IC) device, such as a memory device or buffer device. For example, bond pads, electrostatic discharge  
10   devices, input buffer transistor capacitance, and output driver transistor parasitic and interconnect capacitances relative to the IC device substrate all contribute to the memory device load capacitance.

          The load capacitances connected to multiple points along the length of the signal line may degrade signaling performance. As more load capacitances are introduced along  
15   the signal line of the bus, signal settling time correspondingly increases, reducing the bandwidth of the memory system. In addition, impedance along the signal line may become harder to control or match as more load capacitances are present along the signal line. Mismatched impedance may introduce voltage reflections that cause signal detection errors. Thus, for at least these reasons, increasing the number of loads along  
20   the bus imposes a compromise to the bandwidth of the memory system.

          In an upgradeable memory system, such as conventional memory system 100, different memory capacity configurations become possible. Each different memory capacity configuration may present different electrical characteristics to the

control/address bus 130. For example, load capacitance along each signal line of the control/address bus 130 may change with two different module capacity configurations.

As memory systems incorporate an increasing number of memory module configurations, the verification and validation of the number of permutations that these systems make possible may become increasingly more time consuming. Verification involves the confirmation of operation, logical and/or physical functionality of an IC by running tests on models of the memory, associated devices and/or bus prior to manufacturing the device. Validation involves testing the assembled system or components thereof (e.g., a memory module). Validation typically must account for a majority of the combinations or permutations of system conditions and possibilities which different memory configurations (e.g., 256Mbyte, 1Gbyte...) present including signaling, electrical characteristics (e.g., impedance, capacitance, and inductance variations), temperature effects, different operating frequencies, different vendor interfaces, etc, to name a few. Thus, as the number of possible memory configurations increase, the test and verification time required also increases. More time required to test a system often increases the cost of bringing the system to market or delays a product introduction beyond an acceptable window of time to achieve competitiveness.

There is a need for memory system architectures or interconnect topologies that provide cost effective upgrade capabilities without compromising bandwidth. Using conventional signaling schemes, the bussed approaches lend efficiency towards resource utilization of a system and permits module interfacing for upgradeability. However, the bussed approach may suffer from bandwidth limitations that stem from the electrical characteristics inherent in the bus topology. In addition, impedance along a signal line

may be increasingly more difficult to control with increased connection points along a signal line, introducing impedance mismatch and signal reflections. Utilizing the bussed approach in implementing an upgradeable memory system introduces many possible electrical permutations and combinations with each unique module configuration.

5

## **SUMMARY OF THE INVENTION**

The present invention is directed toward memory system architectures (i.e., interconnect topologies) which include a controller communicating to at least one memory subsystem (e.g., a buffered memory module). An independent point-to-point link may be utilized between the controller and each memory subsystem to eliminate  
10 physical inter-dependence between memory subsystems. According to an embodiment, the memory system may be upgraded by coupling additional memory module(s), each via a dedicated point-to-point link to the controller. Bandwidth may scale upwards as the memory system is upgraded by the additional memory module(s).

In one aspect, the present invention is a memory system comprising a memory  
15 controller having an interface and at least a first memory subsystem. The interface includes a plurality of memory subsystem ports including a first memory subsystem port. The first memory subsystem includes a buffer device having a first port and a second port, and a plurality of memory devices coupled to the buffer device via the second port. A plurality of point-to-point links include a first point-to-point link. Each point-to-point  
20 link has a connection to a respective memory subsystem port of the plurality of memory subsystem ports. The first point-to-point link connecting the first port to a first memory subsystem port to transfer data between the plurality of memory devices and the memory controller.

In another aspect, the present invention is a memory system comprising a controller device and first and second buffer devices, each having a first interface and a second interface. A first point-to-point link includes a first connection to the controller device and a second connection to the first interface of the first buffer device. A first channel is connected to the second interface of the first buffer device, and a first plurality of memory devices are electrically coupled to the first channel. A second point-to-point link includes a first connection to the controller device and a second connection to the first interface of the second buffer. A second channel is connected to the second interface of the second buffer device, and a second plurality of memory devices are electrically coupled to the second channel.

In yet another aspect, the present invention comprises a controller device, and a first and second plurality of buffer devices, each buffer device having a first interface connected to a plurality of memory devices. First and second point-to-point links each include a first end connected to the controller device and a second end connected to a repeater device. A plurality of repeater links couple the first and second repeater devices to respective first and second pluralities of buffer devices.

In another aspect the present invention is a memory system comprising a controller device; a first, second and third connectors; and first second and third point-to-point links. Each of the respective first, second point-to-point links includes a first connection to the interface and a second connection to the respective first, second and third connectors. In this aspect the present invention also includes a first memory subsystem having a buffer device and a plurality of memory devices. The buffer device includes a first interface connected to the first connector, and a second interface



connected to the plurality of memory devices. The second and third connectors may support coupling to respective second and third memory subsystems.

The present invention is described in the detailed description, including the embodiments to follow. The detailed description and embodiments are given by way of illustration only. The scope of the invention is defined by the attached claims. Various modifications to the embodiments of the present invention remain within the scope defined by the attached claims.

### **BRIEF DESCRIPTION OF THE DRAWINGS**

In the course of the detailed description to follow, reference will be made to the attached drawings, in which:

FIGURE 1 illustrates a representational block diagram of a conventional memory system employing memory modules;

FIGURES 2A and 2B illustrate contemporary memory systems employing buffered modules;

FIGURE 3A and 3B illustrate a block diagram representing memory systems according to embodiments of the present invention;

FIGURES 4A, 4B, and 4C illustrate buffered memory modules according to embodiments of the present invention;

FIGURE 5 illustrates a block diagram of a buffer device according to another embodiment of the present invention;

FIGURES 6A and 6B illustrate block diagrams of a memory system according to other embodiments of the present invention;

FIGURE 7 illustrates a block diagram of a memory system employing a buffered quad-channel module according to an embodiment of the present invention;

FIGURE 8A illustrates a block diagram of a large capacity memory system according to another embodiment of the present invention; and

5        FIGURES 8B and 8C illustrate another approach utilized to expand the memory capacity of a memory system in accordance to yet another embodiment of the present invention.

## **DETAILED DESCRIPTION**

The present invention relates to a memory system which includes a plurality of point-to-point links connected to a master. At least one point-to-point link connects at least one memory subsystem to the master, (e.g., a processor or controller). The memory system may be upgraded by coupling memory subsystems to the master via respective dedicated point-to-point links. Each memory subsystem includes a buffer device that communicates to a plurality of memory devices. The master communicates with each buffer device via each point-to-point link. The buffer device may be disposed on a memory module along with the plurality of memory devices and connected to the point-to-point link via a connector. Alternatively, the buffer device may be disposed on a common printed circuit board or backplane link along with the corresponding point-to-point link and master.

“Memory devices” are a common class of integrated circuit devices that have an array of memory cells, such as, dynamic random access memory (DRAM), static random access memory (SRAM), etc. A “memory subsystem” is a plurality of memory devices interconnected with an integrated circuit device (e.g., a buffer device) providing access between the memory devices and an overall system, for example, a computer system. It should be noted that a memory system is distinct from a memory subsystem in that a memory system may include one or more memory subsystems. A “memory module” or simply just “module” denotes a substrate having a plurality of memory devices employed with a connector interface. It follows from these definitions that a memory module having a buffer device isolating data, control, and address signals of the memory devices from the connector interface is a memory subsystem.

With reference to FIGURE 3A and 3B, block diagrams of a memory system according to embodiments of the present invention are illustrated. Memory systems 300 and 305 include a controller 310, a plurality of point-to-point links 320a-320n, and a plurality of memory subsystems 330a-330n. For simplicity, a more detailed embodiment of memory subsystem 330a is illustrated as memory subsystem 340. Buffer device 350 and a plurality of memory devices 360 are disposed on memory subsystem 340. Buffer device 350 is coupled to the plurality of memory devices 360 via channels 370. Interface 375 disposed on controller 310 includes a plurality of memory subsystem ports 378a-378n. A “port” is a portion of an interface that serves a congruent I/O functionality. One of memory subsystem ports 378a-378n includes I/Os, for sending and receiving data, addressing and control information over one of point-to-point links 320a-320n.

According to an embodiment of the present invention, at least one memory subsystem is connected to one memory subsystem port via one point-to-point link. The memory subsystem port is disposed on the memory controller interface which includes a plurality of memory subsystem ports, each having a connection to a point-to-point link.

In FIGURE 3A, point-to-point links 320a-320n, memory subsystems 330a-330c, and controller 310, are incorporated on a common substrate (not shown) such as a wafer or a printed circuit board (PCB) in memory system 300. In an alternate embodiment, memory subsystems are incorporated onto individual substrates (e.g., PCBs) that are incorporated fixedly attached to a single substrate that incorporates point-to-point links 320a-320n and controller 310. In another alternate embodiment illustrated in FIGURE 3B, memory subsystems 330a-330c are incorporated onto individual substrates which include connectors 390a-390c to support upgradeability in memory system 305.

Corresponding mating connectors 380a-380n are connected to a connection point of each point-to-point link 320a-320n. Each of mating connectors 380a-380n interface with connectors 390a-390c to allow removal/inclusion of memory subsystems 330a-330c in memory system 305. In one embodiment, mating connectors 380a-380n are sockets and  
5 connectors 390a-390c are edge connectors disposed on an edge of each substrate 330a-330c. Mating connectors 380a-380n, are attached to a common substrate shared with point-to-point connections 320a-320n and controller 310.

With further reference to FIGURES 3A and 3B, buffer device 350 transceives and provides isolation between signals interfacing to controller 310 and signals interfacing to  
10 the plurality of memory devices 360. In a normal memory read operation, buffer device 350 receives control, and address information from controller 310 via point-to-point link 320a and in response, transmits corresponding signals to one or more, or all of memory devices 360 via channel 370. One or more of memory devices 360 may respond by transmitting data to Buffer device 350 which receives the data via one or more of  
15 channels 370 and in response, transmits corresponding signals to controller 310 via point-to-point link 320a. Controller 310 receives the signals corresponding to the data at corresponding port 378a-378n. In this embodiment, memory subsystems 330a-330n are buffered modules. By way of comparison, buffers disposed on the conventional DIMM module in U.S. Patent 5,513,135 are employed to buffer or register control signals such  
20 as RAS, and CAS, etc., and address signals. Data I/Os of the memory devices disposed on the DIMM are connected directly to the DIMM connector (and ultimately to data lines on an external bus when the DIMM is employed in memory system 100).

Buffer device 350 provides a high degree of system flexibility. New generations of memory devices may be phased in with controller 310 or into memory system 300 by modifying buffer device 350. Backward compatibility with existing generations of memory devices (i.e., memory devices 360) may also be preserved. Similarly, new generations of controllers may be phased in which exploit features of new generations of memory devices while retaining backward compatibility with existing generations of memory devices.

Buffer device 350 effectively reduces the number of loading permutations on the corresponding point-to-point link to one, thus simplifying test procedures. For example, characterization of a point to point link may involve aspects such as transmitters and receivers at opposite ends, few to no impedance discontinuities, and relatively short interconnects. By way of contrast, characterization of control/address bus 130 (see FIGURE 1) may involve aspects such as multiple transmit and receive points, long stub lines, and multiple load configurations, to name a few. Thus, the increased number of electrical permutations tend to add more complexity to the design, test, verification and validation of memory system 100.

Buffered modules added to upgrade memory system 300 (e.g., increase memory capacity) are accommodated by independent point-to-point links. Relative to a bussed approach, system level design, verification and validation considerations are reduced, due to the decreased amount of module inter-dependence provided by the independent point-to-point links. Additionally, the implementation, verification and validation of buffered modules may be performed with less reliance on system level environment factors.

Several embodiments of point-to-point links 320a-320n include a plurality of link architectures, signaling options, clocking options and interconnect types. Embodiments having different link architectures include simultaneous bi-directional links, time-multiplexed bi-directional links and multiple unidirectional links. Voltage or current mode signaling may be employed in any of these link architectures. Clocking methods include any of globally synchronous clocking; source synchronous clocking (i.e., where data is transported alongside the clock) and encoding the data and the clock together. In one embodiment, differential signaling is employed and is transported over differential pair lines. In alternate embodiments, one or more common voltage or current references are employed with respective one or more current/voltage mode level signaling. In yet other embodiments, multi-level signaling-where information is transferred using symbols formed from multiple signal (i.e., voltage/current) levels is employed.

Signaling over point-to-point links 320a-320n may incorporate different modulation methods such as non-return to zero (NRZ), multi-level pulse amplitude modulation (PAM), phase shift keying, delay or time modulation, quadrature amplitude modulation (QAM) and Trellis coding. Other signaling methods and apparatus may be employed in point-to-point links 320a-320n, for example, optical fiber based apparatus and methods.

The term “point-to-point link” denotes one or a plurality of signal lines, each signal line having only two transceiver connection points, each transceiver connection point coupled to transmitter circuitry, receiver circuitry or transceiver circuitry. For example, a point-to-point link may include a transmitter coupled at or near one end and a

receiver coupled at or near the other end. The point-to-point link may be synonymous and interchangeable with a point-to-point connection or a point-to-point coupling.

In keeping with the above description, the number of transceiver points along a signal line distinguishes between a point-to-point link and a bus. According to the above,  
5 the point-to-point link consists of two transceiver connection points while a bus consists of more than two transceiver points.

One or more terminators (e.g., a resistive element) may terminate each signal line in point-to-point links 320a-320n. In several embodiments of the present invention, the terminators are connected to the point-to-point link and situated on buffer device 350, on  
10 a memory module substrate and optionally on controller 310 at memory subsystem ports 378a-378n. The terminator(s) connect to a termination voltage, such as ground or a reference voltage. The terminator may be matched to the impedance of each transmission line in point-to-point links 320a-320n, to help reduce voltage reflections.

In an embodiment of the present invention employing multi-level PAM signaling,  
15 the data rate may be increased without increasing either the system clock frequency or the number of signal lines by employing multiple voltage levels to encode unique sets of consecutive digital values or symbols. That is, each unique combination of consecutive digital symbols may be assigned to a unique voltage level, or pattern of voltage levels. For example, a 4-level PAM scheme may employ four distinct voltage ranges to  
20 distinguish between a pair of consecutive digital values or symbols such as 00, 01, 10 and 11. Here, each voltage range would correspond to one of the unique pairs of consecutive symbols.



With reference to FIGURES 4A, 4B and 4C, buffered memory modules according to embodiments of the present invention are shown. Modules 400 and 405 include buffer device 405 and a plurality of memory devices 410a-410h communicating over a pair of channels 415a and 415b. In these embodiments channel 415a communicates to memory  
5 devices 410a-410d and channel 415b communicates to memory devices 410e-410h.

In an embodiment, channels 415a and 415b consist of a plurality of signal lines in a relatively short multi-drop bus implementation. The plurality of signal lines may be controlled impedance transmission lines that are terminated using respective termination elements 420a and 420b. Channels 415a and 415b are relatively short (i.e., are coupled  
10 to relatively few memory devices relative to a conventional memory system, for example see FIGURES 2A and 2B) and connect to an I/O interface (not shown) of each memory device via a short stub. Signal lines of channels 415a and 415b include control lines (RQ), data lines (DQ) and clock lines (CFM, CTM). The varieties of interconnect topologies, interconnect types, clocking methods, signaling references, signaling  
15 methods, and signaling apparatus described above in reference to point-to-point links 320a-320n may equally apply to channels 415a and 415b.

In accordance with an embodiment of the present invention, control lines (RQ) transport control (e.g., read, write, precharge...) information and address (e.g., row and column) information contained in packets. By bundling control and address information  
20 in packets, protocols required to communicate to memory devices 410a-410h are independent of the physical control/address interface implementation.

In alternate embodiments, control lines (RQ) may comprise individual control lines, for example, row address strobe, column address strobe, etc., and address lines.

Individual point-to-point control and address lines increase the number of parallel signal connection paths, thereby increasing system layout resource requirements with respect to a narrow “packet protocol” approach. In one alternate embodiment illustrated in FIGURE 6A, individual device select lines 633a and 633b are employed to perform  
5 device selection. Individual device select lines 633a and 633b decrease some latency consumed by decoding device identification which normally is utilized when multiple devices share the same channel and incorporate individual device identification values.

Clock lines of channels 415a and 415b include a terminated clock-to-master (CTM) (i.e., clock to buffer) and clock-from-master (CFM) (i.e., clock from buffer) line.  
10 In a source synchronous clocking method, CTM may be transition or edge aligned with control and/or data communicated to buffer device 405 from one or more of memory devices 410a-410d in, for example, a read operation. CFM may be aligned with or used to synchronize control and/or data from the memory buffer to buffer device 405 in, for example, a write operation.

15 Although two channels 415a and 415b are shown in FIGURE 4A, a single channel is also feasible. In other embodiments, more than two channels may be incorporated onto module 400. It is conceivable that if each channel and memory device interface is made narrow enough, then a dedicated channel between each memory device and the buffer device may be implemented on the module. The width of the channel  
20 refers to the number of parallel signal paths included in each channel. FIGURE 4B illustrates a quad-channel module 450 having channels 415a-415d. In this embodiment, channels 415c and 415d are routed in parallel with channels 415a and 415b to support more memory devices (e.g., 32 memory devices). By incorporating more channels and

additional memory devices, module 400 (FIGURE 4B) may be implemented in memory systems that require large memory capacity, for example, in server or workstation class systems.

5 In alternate embodiments, channels 415a and 415b may operate simultaneously with channels 415c and 415d to realize greater bandwidth. By operating a plurality of channels in parallel, the bandwidth of the module may be increased independently of the memory capacity. The advantages of greater bandwidth may be realized in conjunction with larger capacity as more modules incorporated the memory system 305 (see FIGURE 3B) increase the system memory capacity. In other alternate embodiments, the modules  
10 are double sided and channels along with corresponding pluralities of memory devices are implemented on both sides. Using both sides of the module increases capacity or increases bandwidth without impacting module height. Both capacity and bandwidth may increase using this approach. Indeed, these techniques may increase capacity and bandwidth singly or in combination.

15 Other features may also be incorporated to enhance module 400 in high capacity memory systems, for example, additional memory devices and interface signals for error correction code storage and transport (ECC). Referring to FIGURE 4C, memory devices 410i and 410r intended for ECC are disposed on module 470.

In one embodiment, memory devices 410a-410h are Rambus Dynamic Random  
20 access Memory (RDRAM) devices operating at a data rate of 1066Mbits/sec. Other memory devices may be implemented on module 400, for example, Double Data Rate (DDR) DRAM devices and Synchronous DRAM (SDRAM) devices. Utilizing buffer device 405 between the memory devices and controller in accordance with the present

invention (e.g., see FIGURE 3) may feasibly render the type of memory device transparent to the system. Different types of memory devices may be included on different modules within a memory system, by employing buffer device 405 to translate protocols employed by controller 310 to the protocol utilized in a particular memory device implementation.

With reference to FIGURE 5, a block diagram of a buffer device according to an embodiment of the present invention is illustrated. Buffer device 405 includes interface 510, interface 520a and 520b, multiplexing 530a and 530b, control logic 540, write buffer 550, optional cache 560, computation block 565, clock circuitry 570 and operations circuitry 572.

In an embodiment, interface 510 couples to external point-to-point link 320 (e.g., point-to-point links 320a-320n in FIGURES 3A and 3B). Interface 510 includes a port having transceiver 575 (i.e. transmit and receive circuitry) that connects to a point-to-point link. Point-to-point link 320 comprises one or a plurality of signal lines, each signal line having no more than two transceiver connection points. One of the two transceiver connection points is included on interface 510. Buffer device 405 may include additional ports to couple additional point-to-point links between buffer device 405 and other buffer devices on other memory modules. These additional ports may be employed to expand memory capacity as is described in more detail below. In the embodiment shown in FIGURE 5, buffer device 405 may function as a transceiver between point-to-point link 320 and other point-to-point links.

In one embodiment, termination 580 is disposed on buffer device 405 and is connected to transceiver 575 and point-to-point link 320. In this embodiment, transceiver

575 includes an output driver and a receiver. Termination 580 may dissipate signal energy reflected (i.e., a voltage reflection) from transceiver 575. Termination 580 may be a resistor or capacitor or inductor, singly or a series/parallel combination thereof. In alternate embodiments, termination 580 may be external to buffer device 405. For example, termination 580 may be disposed on a module substrate or on a memory system substrate.

In another approach, signal energy reflected from transceiver 575 may be utilized in a constructive manner according to an embodiment. By correctly placing a receive point spaced by a distance from the end of point-to-point link 320, a reflected waveform is summed with an incident waveform to achieve a greater signal amplitude. In this approach, layout space may be saved by eliminating termination 580. System power may also be saved using this approach since smaller incident voltage amplitude waveforms may be employed. This approach may be equally applicable to the transceiver end of the point-to-point link, or to channels 415a and 415b (see FIGURES 4A to 4C).

With further reference to FIGURE 5, interfaces 520a and 520b receive and transmit to memory devices disposed on the module (e.g., see FIGURES 4A, 4B and 4C) via channels. Ports included on interfaces 520a and 520b connect to each channel. In alternate embodiments of the present invention, interfaces 520a and 520b include any number of channels e.g., two, four, eight or more channels.

According to an embodiment of the present invention, multiplexers 530a and 530b perform bandwidth-concentrating operations, between interface 510 and interfaces 520a and 520b. The concept of bandwidth concentration involves combining the (smaller) bandwidth of each channel in a multiple channel embodiment to match the

(higher) overall bandwidth utilized in a smaller group of channels. This approach typically utilizes multiplexing and demultiplexing of throughput between the multiple channels and smaller group of channels. In an embodiment, buffer device 405 utilizes the combined bandwidth of interfaces 520a and 520b to match the bandwidth of interface

5 510. Bandwidth concentration is described in more detail below.

Cache 560 is one performance enhancing feature that may be incorporated onto buffer device 405. Employing a cache 560 may improve memory access time by providing storage of most frequently referenced data and associated tag addresses with lower access latency characteristics than those of the memory devices. Computation

10 block 565 may include a processor or controller unit, a compression/decompression engine, etc, to further enhance the performance and/or functionality of the buffer device. In an embodiment, write buffer 550 may improve interfacing efficiency by utilizing available data transport windows over point-to-point link 320 to receive write data and optional address/mask information. Once received, this information is temporarily stored

15 in write buffer 550 until it is ready to be transferred to at least one memory device over interfaces 520a and 520b.

A serial interface 574 may be employed to couple signals utilized in initialization of module or memory device identification values, test function, set/reset, access latency values, vendor specific functions or calibration. Operations circuitry 572 may include

20 registers or a read-only memory (ROM) to store special information (e.g., vendor or configuration information) that may be used by the controller. Operations circuitry may reduce costs by eliminating the need for separate devices on the module conventionally

provided to perform these features (e.g., serial presence detect (SPD) employed in some conventional DIMM modules).

According to an embodiment of the present invention, sideband signals are employed to handle special functions such as reset, initialization and power management  
5 functions. Sideband signals are connected via serial interface 574 and are independent from point-to-point link 320 for handling the special functions. In other embodiments sideband signals are independently coupled to memory devices 410a-410h to directly promote initialization, reset, power-up or other functionality independently of buffer device 405. Other interconnect topologies of sideband signals are possible. For example,  
10 sideband signals may be daisy chained between buffer devices and coupled to the memory controller or daisy chained between all memory devices to the memory controller. Alternatively, dedicated sideband signals may be employed throughout.

Clock circuitry 570 may include clock generator circuitry (e.g., Direct Rambus Clock Generator) which may be incorporated onto buffer device 405 and thus may  
15 eliminate the need for a separate clock generating device. Here, module or system costs may be decreased since the need for a unique clock generator device on the module or in the system may be eliminated. Since reliability to provide adequate clocking on an external device is eliminated, complexity is reduced since the clock may be generated on the buffer device 570. By way of comparison, some of the conventional DIMM modules  
20 require a phase lock loop (PLL) generator device to generate phase aligned clock signals for each memory device disposed on the module.

According to an embodiment of the present invention, clocking circuitry 570 includes one or more clock alignment circuits for phase or delay adjusting internal clock

signals with respect to an external clock (not shown). Clock alignment circuitry may utilize an external clock from an existing clock generator, or an internal clock generator to provide an internal clock, to generate internal synchronizing clock signals having a predetermined temporal relationship.

5           With reference to FIGURES 6A, and 6B, block diagrams of a memory system according to embodiments of the present invention are illustrated. Memory system 600 includes modules 400a and 400b, controller 610, and populated primary point-to-point links 620a and 620b. Unpopulated primary point-to-point links 630 are populated by coupling additional modules (not shown) thereto. The additional modules may be  
10   provided to upgrade memory system 600. Connectors may be disposed at an end of each primary point-to-point link to allow insertion or removal of the additional modules. Modules 400a and 400b may also be provided with a connector or may be fixedly disposed (i.e., soldered) in memory system 600. Although only two populated primary point-to-point links are shown in FIGURE 6A, any number of primary point-to-point  
15   links may be disposed in memory system 600, for example, three primary point-to-point links 400a-400c, as shown in FIGURE 6B.

          With reference to FIGURE 7 and FIGURE 4B, a block diagram of a memory system employing a buffered quad-channel module according to an embodiment of the present invention is illustrated. Memory systems 700 incorporate quad-channel modules  
20   450a-450d, each coupled via point-to-point links 620a-620d respectively.

          Referring to FIGURE 4B, buffer device 405 may operate in a bandwidth concentrator approach. By employing quad channels 415a-415d on each of modules 450a-450d, bandwidth in each module may be concentrated from all quad channels 415a-



415d on each module to corresponding point-to-point links 620a-620d. In this embodiment, throughput on each of point-to-point links 620a-620d is concentrated to four times the throughput achieved on each of quad channels 415a-415d. Here, each of channels 415a-415d transfers information between one or more respective memory  
5 devices on each channel and buffer device 405 simultaneously.

Any number of channels 415a-415d, for example; two channels 415c and 415d may transfer information simultaneously and the memory devices on the other two channels 415a and 415b remain in a ready or standby state until called upon to perform memory access operations. Different applications may have different processing  
10 throughput requirements. In addition, the throughput requirements of a particular application may dynamically change during processing. Typically, more power is consumed as throughput is increased as power consumption relates in proportion to operation frequency. The amount of throughput in a system may be implemented on a dynamic throughput requirement basis to save on power consumption. In this  
15 embodiment, memory system 700 may concentrate bandwidth as it is required while in operation. For example, memory system 700 may employ only one of channels 415a-415d and match throughput to the corresponding point-to-point link. As bandwidth requirements increase, memory system 700 may dynamically activate more of channels 415a-415d and increase the throughput on the point-to-point link along with the number  
20 of channels accordingly to meet the bandwidth requirements for a given operation.

With reference to FIGURE 8A, a block diagram of a large capacity memory system according to an embodiment of the present invention is illustrated. Memory system 900 includes modules 470a-470p, coupled to controller 610 via repeaters 910a-

910d, primary links 920a-920d, and repeater links 930a-930p. Primary links 920a-920d provide a point to point link between controller 610 and a respective repeater 910a-910d. In an embodiment of the present invention, each of repeaters 910a-910d decode packets transmitted from controller 610 which are then directed over one or more, or none of  
5 repeater links 930a-d, depending the type of access required. Each repeater link 930a-930p may utilize a point-to-point link configuration. By incorporating, repeated links 930a-930p and repeaters 910a-910d, a larger number of modules may be accessed and a larger capacity memory system may be realized. Such a large capacity may be suited in a computer server system.

10       FIGURE 8B illustrates another approach utilized to expand the memory capacity of a memory system in accordance to yet another embodiment. Here, a plurality of buffered modules 950a-950d are “daisy chained” via a plurality of point-to-point links 960a-960d to increase the overall memory capacity. Connection points of each point-to-point link are connected to two adjacent buffered modules. Each of buffered modules  
15 950a-950c transceive signals between adjacent point-to-point links 960a-960d. Point-to-point link 960a may be coupled to a controller or another buffered module. Additional point-to-point links may be coupled to a buffer device in a tree configuration approach. For example, three point-to-point links 970a-970c each having a single end connected to one buffer device may be employed as shown in FIGURE 8C.

20       While this invention has been described in conjunction with what is presently considered the most practical embodiments, the invention is not limited to the disclosed embodiments. In the contrary, the embodiments disclosed cover various modifications that are within the scope of the invention as set forth in the following claims.